First draft of paper on Jaynes' proof.
Authors: P. Castell and Prof. M. Redhead.
Cambridge, February 1990.

Jaynes' Subjective H-Theorem

Introduction

One of the most fundamental problems in trying to reduce thermodynamics to statistical mechanics is to give an account of irreversibility. It is generally thought that the Gibbs fine-grained entropy cannot do this because it remains constant as a consequence of Liouville's theorem. Gibbs famously resorted to coarse-graining, but many critics of the Gibbs approach claim that coarse-graining an artificial mathematical procedure which has no physical justification. In 1965 E.T. Jaynes published a paper in which he claims to have discovered an almost unbelievably simple proof of the second law of thermodynamics. Our aim is to show in greater detail how the proof works and to correct some widespread misunderstandings about its significance. In our opinion this proof is a major contribution in the effort to clarify the foundations of statistical physics.

Although Jaynes calls the proof the 'subjective H-theorem', it does not explicitly depend on any particular philosophical interpretation of probability. One of the great advantages of the subjective approach to statistical mechanics over its objectivist rivals is that it gives us relatively easy access to probabilities. The maximum entropy formalism makes no microscopic assumptions (e.g. about molecular collisions) and does not invoke the ergodic hypothesis. This does, however, mean that two main approaches to proving the second law have been ruled out. This need not worry the subjectivists since these approaches have not been altogether successful, but it is of great importance that they provide an alternative approach to the second law. The subjective H-theorem is a prime candidate for such an alternative.

Part of what makes the subjective H-theorem interesting and confusing is a conceptual worry connected with the subjective interpretation of entropy: prima facie, at least, it is not at all clear why increases in experimental entropy should have anything to do with the subjective entropy analog, which, on the Jaynes interpretation, just reflects our degree of ignorance about the true state of the system. This question will be addressed at the end of the paper.

A note on terminology

The subjective H-theorem is needed to fill out the maximum entropy story, but this does not mean that the proof cannot be used by objectivists. As a result there are some choices to be made concerning our terminology. For instance, when we say that the canonical distribution 'appropriately describes a system' in statistical equilibrium, we are not necessarily implying that the canonical distribution is a property of that system. We shall generally use the subjectivist terminology when discussing the proof, but shall provide an objectivist rephrasing of the crucial points in one section of the paper.

Since the proof crucially depends on the distinction between two senses of the word 'equilibrium' a note on our terminology is in order. We will say that a system is in *macroscopic equilibrium* when the macroscopic parameters that define the thermodynamic state of the system are not changing over time. A macroscopic parameter is considered to have reached an equilibrium value when its measured fluctuations are of the same order as the accuracy of the readings.

The second sense of 'equilibrium' is *statistical equilibrium.* A system in macroscopic equilibrium is said to be in statistical equilibrium when it is appropriately described by the canonical distribution.

The proof

The proof relies on the following facts:

i) The Gibbs fine-gained entropy S_G stays constant in an adiabatic system.

ii) S_G is related to the experimental entropy S_e by the inequality $S_G \leq S_e$.

iii) S_G is equal to S_e if and only if S_G is evaluated over the canonical distribution.

iv) The canonical distribution appropriately describes a system in statistical equilibrium.

The scenario is as follows:

1) Consider a system in *statistical* equilibrium at time t_1. It is contained by a rigid diathermic partition which is positioned in a heat bath. The average energy of the canonical ensemble is denoted by $\langle E \rangle$.
2) We now adiabatically isolate the system so that there is no heat exchange when the system is out of macroscopic equilibrium.
3) We then do some irreversible work on the system: S_G stays constant, while $\langle E \rangle \rightarrow \langle E \rangle'$.
4) We leave the system to achieve *macroscopic* equilibrium at time t_2.
5) We remove the isolation and let the system 'relax' to *statistical* equilibrium.

The proof:

a) At time t_1: $S_G = S_e$

b) At time t_2: $S_G \leq S_e'$

c) S_G has remained constant, therefore $S_e \leq S_e'$

Entropy as a dispositional quality

Thermodynamic entropy (S_e) is a very curious concept. It is similar to other state variables such as temperature and pressure in the sense that it is a quantity which describes the macroscopic state of a system in equilibrium. Entropy is, however, defined by the differential equation

$$dS_e = \left.\frac{\delta Q}{T}\right)_{rev} \qquad [I]$$

where δQ is the amount of heat reversibly absorbed by the system at temperature T. Entropy changes are thus given by the integral

$$(S_e)_2 - (S_e)_1 = \int_1^2 \left.\frac{\delta Q}{T}\right)_{rev} \qquad [II]$$

Note that the entropy change between states 1 and 2 involves an integral over a reversible path, i.e., the system is effectively in a succession of equilibrium states.

Now consider the scenario described above. From the thermodynamic point of view one wants to say that $S_e \rightarrow S_e'$ during stage (3). But notice that no reversible change has taken place here, so it is not immediately clear that we can talk about an increase in entropy at all. Statements about entropy changes in irreversible processes must be seen as counterfactuals. They say that *if we had* got from state 1 to state 2 via a reversible path *we would have* got a certain entropy increase as defined by equation [II].

Entropy can therefore be thought of as a dispositional property (see Denbigh and Denbigh (1985), p.48). We might say that one glass is more fragile than another because *if* we dropped them both under a specified set of conditions one *would* break but the other *would* not. Similarly, we say that state 1 of a system has lower entropy than state 2 because *if* we tried to get from state 1 to state 2

via any reversible path there *would* be heat exchange between the system and the surroundings which *would* lead the integral in equation [II] to converge to the same finite positive value.

The crucial point here is that when we give a thermodynamic gloss of the increase of the thermodynamic entropy $S_e \rightarrow S_e'$ (stage (3) of the scenario) we must actually talk about what would have happened to the gas if it had got from state 1 to state 2 via a reversible path. We must therefore talk about a scenario in which the gas is not isolated from its surroundings.

Possible objections to the proof

Objection α) At stage 1 one assumes that the canonical distribution is the appropriate description of the system. One might, however, object that the canonical distribution is no longer appropriate at stage 2, since the heat bath has effectively been removed. The interaction forces between the molecules of the system and those of the heat bath will be broken by the isolation process. It is precisely these forces which allowed small heat fluctuations across the diathermic partition in stage 1. If the isolation process occurs instantaneously the actual state of the system will not change, but the long run behaviour of the system will change because the energy of the system is now constrained to a definite value, whereas before only the long run average energy was fixed. Removing the heat bath will not affect any of the measurable macroscopic parameters, but it will affect the evolution of the microstate of the system.

Thus one might argue that the canonical distribution is no longer appropriate at stage 2 and that one ought to employ the microcanonical distribution instead. Ex hypothesi we know that the energy is fixed at stage 2, but unfortunately we can only guess at its value. The 'best guess' is naturally $\langle E \rangle$ as defined by the canonical distribution, but to assume that this is actually the true energy of the system at stage 2 would be to take information that is not available to us. Our conclusion then, is that the canonical distribution is still appropriate to the system at stage 2.

Objection β) Having denied that the distribution which appropriately describes the system changes as we move from stage 1 to stage 2, the same question arises with regard to the stage $4 \rightarrow 5$ transition. It is not immediately clear why the non-canonical distribution of stage 4 should relax to the canonical distribution in stage 5. There seems to be a curious asymmetry: our distribution changes when we introduce the heat bath, but does not change when we remove it. How does this asymmetry come about? ?

The canonical distribution of stage 2 time-developes into a non-canonical distribution as irreversible work is performed on the adiabatic system. The hamiltonian of the system is modified by a time-dependent factor which is responsible for the irreversible change. But when we perform the stage $4 \rightarrow 5$ transition the hamiltonian which controls the development of the system is no longer just that of the system, it is that of the system and the heat bath conjoined. The original system is now just a subsystem and as such will approach the canonical distribution. There is a crucial difference between the $1 \rightarrow 2$ transition and the $4 \rightarrow 5$ transition. In the former case we do not know the exact energy of the isolated system but we do know the expectation value of its energy. Thus we are justified in applying the canonical distribution. In the latter case we start off with a non-canonical distribution and then introduce a heat bath. The time-development of the system is no longer predictable using Liouville's equation because we do not know the hamiltonian of the heat bath which affects the evolution of the system. We cannot prove that the distribution will relax to the canonical one because we no longer know the relevant hamiltonian. Given this lack of knowledge we assume that the system will return to statistical equilibrium after having been in macroscopic equilibrium for a certain period.

Objection γ) A third possible objection is that the increase of the Gibbs entropy S_G only occurs at stage 5 of the scenario. According to the text-book tradition of thermodynamics one would, however, expect the increase of S_e to occur during stage 3. The worry is that the statistical entropy analog S_G does not change at the same time as its supposed thermodynamic counterpart S_e.

An answer to this objection has already been indicated. It is just to deny that S_e changes to S_e' during the irreversible process. There is no way to determine the experimental entropy during stage 3 of the scenario because the system is not even in macroscopic equilibrium. But is it possible to determine S_e for stage 4? The hypothetical experiment that defines what we mean by the change in S_e must leave the system in contact with the heat bath throughout the change.

One could simply try to define the experimental entropy of state 1 to be the same as state 2, and that of state 4 to be the same as state 5. A putative justification for such a definition might be that none of the other macroscopic variables change detectably between these state pairs. Now entropy can be defined in terms of these other variables using Gibbs' equation and so it ought not to change either. In this way one could determine S_e for stage 4.

Such an account is pleasing in the sense that it satisfies our intuition about where the experimental entropy change takes place, namely during stage 3. The consequences of such an account are that an increase in S_e is not be accompanied by a simultaneous increase in S_G. But if the experimental entropy must be defined in terms of a hypothetical experiment we could (and maybe should) consider the Gibbs entropy increase for that hypothetical setup. Here the distribution describing the system will always be canonical and as such $S_G = S_e$ throughout.

Objection δ) The whole proof rests on the fact that if you start with the canonical distribution and change the energy of the system adiabatically, then you must initially diverge from that distribution. If you start with a non-canonical distribution and then and let the energy of the system fluctuate by introducing a heat bath, then you may approach the canonical distribution. Like Gibbs' generalised H-theorem the subjective H-theorem does not show that any *monotonic increase* in the entropy will occur. Nor is there an indication of how much the distribution will diverge from the canonical one.

Jaynes' proof assumes that we can characterise the initial state of the system by the canonical distribution, but he himself admits that we must monitor the past (macroscopic) behaviour of the system "sufficiently long" in order to tell whether it is in statistical equilibrium, and that "only experience can tell the experimenter how long is "sufficiently long"".

Our worry is that the proof is too dependent on the initial distribution. In fact it crucially depends on starting with precisely the canonical distribution. If we however choose a non-canonical distribution which is close enough to the canonical one to make all the required macroscopic predictions, then the proof no longer goes through. This objection naturally does not invalidate the proof, but it illustrates the use of a theorem which shows that the proof has some resiliency in the face of non-canonical initial distributions.

Reproducibility and entropy increase

The concept of reproducibility plays an important role in our understanding of irreversibility. All thermodynamic experiments are assumed to be reproducible in the sense that a given macrostate evolution depends only on the initial macroscopic state and the macrosopic parameters that control the evolution of the system. In a reproducible system two distinct initial macro-states may develop into the same final macrostate, but one cannot allow that one initial macro-state sometimes develops into one macro-state and sometimes into another under the same macroscopic conditions. Thermodynamic processes are reproducible and will therefore usually lead to certain macroscopic 'equilibrium states'. These states are just those macro-states which comprise the largest number of micro-states.

The statistical theory predicts a probability distribution over micostates that evolves from the initial probability distribution by Liouville's equation. The 'high-probability states' of the initial distribution will always be included in the set of 'high-probability states' of the predicted distribution

if the process is to be reproducible (and predicts the right results). In addition there may well be other 'high-probability states' in the predicted distribution. Thus the very fact that thermodynamic processes are reproducible implies that the number of 'high-probability states' must always get larger.
On the original Boltzmann interpretation this means that the entropy cannot decrease.

The leading idea behind the Jaynes proof is that S_G only stays constant in the given scenario because one approaches the high entropy macrostate from one particular initial state. In contrast "the experimental entropy is a measure of all conceivable ways in which the final macrostate can be realised, and not merely all the ways in which it could be produced in one particular experiment" (Jaynes, 1965).

The scenario underlying the subjective H-theorem shows us that we can gain knowledge of the micro-state of a system from a knowledge of its macroscopic history. Macro-states are essentially collections of micro-states and the experimental entropy of a macro-state is proportional to the logarithm of the number of microstates that comprise that state. The asymmetry comes about because knowledge of past macro-states not only lets you predict future ones, but also tells you which subclass of micro-states the system will belong to. In fact Liouville's theorem tells us that our degree of ignorance with respect to the actual micro-state of the system will not change during the adiabatic evolution. Using the Liouville equation we can conserve our information about the true micro-state and hence the subjective Gibbs entropy stays constant.

The experimental entropy must always be greater or equal to the subjective entropy, *given the same constraints on each*, but if there are important physical constraints of which the experimenter is not aware, the experimental entropy will be lower than the predicted value, so that the inequality $S_G \leq S_e$ no longer holds in general. It only holds for the case when all the relevant physical constraints are known (and imposed on the subjective probability distribution).

In the scenario above, one assumes that $S_G = S_e$ at stage 1, and therefore we are dealing with the special case in which all relevant constraints have been taken into account. One also assumes that the work done on the system in stage 3 is known. Because the system's evolution is assumed to be a macroscopically reproducible one, the system is bound to finish in a state of higher experimental entropy.

On the orthodox interpretation of statistical mechanics, the constancy of the Gibbs entropy is is hard to explain. The change of experimental entropy is evidently not mirrored by an equivalent change in the Gibbs analog. But if one has made a correct assessment of the true objective probability distribution in stage 1 and then changed it according to Liouville's equation, then surely that new distribution ought to tell you the experimental entropy in stage 5. The subjectivist, on the other hand, primarily regards the Gibbs entropy as a measure of our ignorance as to the true micro-state. This does not change so long as no unknown constraints are added without our knowledge.

But, granted that the subjectivist is not trying to mirror the change in experimental entropy with the subjective analog, how can he use this analog to demonstrate the second law? There is a link between the experimental entropy and the subjective analog: The experimental entropy corresponds to that subjective entropy which represents a state of knowledge which is confined simply to the physical macroscopic constraints. (This does not include knowledge about constraints which were previously imposed on the system). The experimental (thermodynamic) entropy is what Harold Grad (1967) calls appropriate to a low level description.

If we simply choose to deal with those macro-states that behave reproducibly then it is not surprising that they show entropic irreversibility on a phenomelogical level. Indeed, Jaynes himself says (1988) that he does not see "why anymore than [reproducibility] is needed to understand and explain the observed phenomenological irreversibility of thermodynamics".

Reproducibility enters the Jaynes' proof when one assumes that the time-developed distribution function leads to a correct prediction of the thermodynamic variables that define the new macroscopic state. As we shall see in the last section it is the central objective criterion on which the whole proof hinges.

Framing and coarse-graining

Hollinger and Zenzen (1982, pp.327-328) consider the Jaynes proof to demonstrate the *a priori* irreversible nature of probabilistic prediction. They claim that the proof cannot, however, explain "the entropic irreversibility of observed phenomena". As in coarse-graining, irreversibility is seen as an artificial consequence of the mathematical procedure. Indeed, Hollinger and Zenzen believe that the a process they call "framing" underlies the Jaynes proof and is simply coarse-graining in its "most primitive form".

Gibbs' generalised H-theorem depends on a spreading of weighted points throughout phase space. This spreading effect is measured by the coarse-grained H-function and the function is shown to decrease. The theorem relies on the fact that the highest possible value of H is obtained when all the weighted points are contained in exactly one of the cells into which the space is divided. One assumes that the system begins in such a state and so it is not surprising that the H-function cannot increase. The weighted points cannot contract into one part of the cell they started off in because Liouville's theorem tells us that points in phase-space behave like an incompressible fluid. But, even if such a contraction could take place the coarse-grained entropy could not fall, because it is calculated by averaging the weight density over the cell.

Hollinger and Zenzen do not explicitly state whether they think that 'framing' is part of what Jaynes describes in his 1965 paper, but they certainly leave room for such a misinterpretation, if that's what it is. We wish to point out that 'framing' (considered as a special case of coarse-graining) is fundamentally different from the argument underlying the Jaynes proof.

In the Jaynes proof one assumes that not only the initial macrostate is given, but also that the system is appropriatly described by the canonical distribution. In this case the entropy of the final distribution can never be larger than the entropy of the initial distribution. The canonical distribution encodes the state of knowledge corresponding to the macro-state, but nothing more. The distribution will typically evolve to a non-canonical one. The new macroscopic state will, however, be describable by a new canonical distribution corresponding to the new energy. Therefore the time-evolved distribution will necessarily encode more information than the predicted macro-state.

The irreversibility demonstrated by coarse-graining is entirely artificial in the sense that it would occur regardless of the actual behaviour of the system, or indeed an imaginary ensemble of systems. The subjectrive H-theorem, on the other hand, demonstrates the increase of experimental entropy in a fundamentally different way. It does not try to show that a *statistical entropy analog* must always increase, instead it relies on the combination of two facts; firstly that the Gibbs fine-grained analog stays constant and secondly that the macro-state evolution is reproducible. (The proof assumes that Liouvilles equation will predict the correct macrostate). It is not the *change of shape* of the weighted region in phase space that is supposed to be responsible for the entropy increase, rather, it is the *translation* of the 'high probability region' in phase space together with Liouville's theorem that forces the distribution to become non-canonical. The new distribution encodes information about the history of the system aswell as its new macro-state.

Does the subjective H-theorem show there to be a truely physical irreversibility, or does it just tell us something about the nature of probabilistic prediction as Hollinger and Zenzen maintain? The answer is that it does demonstrate physical irreversibility, but only on the assumption of reproduciblity.

To many, the assumption of reproducibility will seem far too strong. Instead of a demonstration based on the underlying laws of particle motion we have only shown that macroscopically reproducible processes will result in an increase of thermodynamic entropy. The question which will be asked now is what features of a process make it a reproducible one.

Ergodic theory attempts to explain why Gibbs phase averaging works and to justify the use of the canonical and microcanonical distributions. The subjective H-theorem does not explain these things. Subjectivists such as Jaynes tell an additional story about ignorance priors to answer these questions.

Subjectivity and reduction

Traditionally the relation between thermodynamics and statistical physics has always been seen as a paradigm case of reduction. Thermodynamic variables such as temperature and pressure have found convincing new interpretations in terms of a statistical model[1]. On the objective interpretation of Gibbs' statistical mechanics all these macroscopic quantities are not to be seen as properties of the system per se, instead they are to be seen as properties of an ensemble of systems, all of which are relevantly similar to the actual system. The ergodic theorem is supposed to tell us what ensemble averages have to do with properties of the actual system. Unfortunately it only tells us that they equal *infinite* time-averages and so the story needs an additional theorem to tell us what infinite time-averages have to do with the finite time-averages of measurement.

The subjective interpretation of the Gibbs method does not face the same foundational difficulties, (although it faces others). On this interpretation an ensemble average *is* simply our best estimate of the actual quantity because the probability distribution over which we average represents our degrees of belief about which member of the ensemble is actual. On the subjectivist interpretation the Gibbs entropy analog explicitly becomes a function of what we believe about the system, rather than a property of the system itself. In Jaynes' words entropy "measures the extent of human ignorance as to the microstate" (Jaynes 1965). Unlike temperature and pressure which depend on the *expected* energy of the system, entropy only seems to depend on how peaked our probability distribution for the energy is. On the Jaynes interpretation the entropy is a function of the available evidence about the system. On the objectivist interpretation only the *predicted value* of the entropy is a function of the evidence.

The most obvious objection to the subjectivist account is that if God told you the exact microstate of the system the thermodynamic entropy would not change one bit, but the subjectivist analog would go to zero[2]. There are two ways in which one could try to evade this problem. The first would be to insist that we can only gain information about a system by imposing macroscopic constraints which would actually lower the thermodynamic entropy (e.g. Brillouin, 1962). In this way our subjective entropy could never be lower than the thermodynamic entropy.

But even if this is allowed it may be possible that we have less information available than is necessary to correctly predict the thermodynamic entropy. If the system is cooled down, the thermodynamic entropy will fall regardless of whether you know of the temperature change. But on the subjective interpretation the Gibbs entropy analog will not change. Jaynes says that the maximum entropy formalism will only produces sharp predictions when there is sufficient information to do so, since it always chooses the most 'modest' probability distribution. But in the case of entropy itself this is evidently not the case since we can deduce the Gibbs entropy from the probability distribution itself.

The prediction of experimental entropy can, therefore, go drastically wrong even if no false assumptions have been made. S_e is dependent on the number of macroscopic constraints imposed on the system. A subjective analog of S_e will have to involve an estimate of the degree to which the system is constrained. The Jaynes (maximum entropy) analog makes the most modest estimate, i.e. it assumes there are no other constraints on the system apart from those that are 'known' to exist [3]. This means that the Jaynes analog will usually *underestimate* the experimental entropy (rather than overestimate it). One should, however, not see this in contradiction to the inequality $S_G \leq S_e$. This inequality only compares the subjective entropies of probability distributions which are compatible with *the same* macroscopic constraints.

[1] The case of entropy is not, however, so straightforward. Various intuitive notions of disorder can be found throughout the literature, but none of these has been made precise in a satisfactory way, see for instance Denbigh (1989).

[2] Of course, on the objectivist account, probabilities are independent of your state of knowledge and so the Gibbs entropy analog is too.

[3] 'Known' in this context indicates that the experimenters have a high degree of belief in the presence of those constraints.

The second option, and this is the one Jaynes takes, is to deny that thermodynamic entropy is itself objective. We quote: "*Even at the purely phenomenological level, entropy is an anthropomorphic concept*.. For it is a property, not of the physical system, but of the particular experiments you or I choose to perform on it" (Jaynes (1965)). Following Denbigh and Denbigh (1985) we deny that thermodynamic entropy is a subjective quantity under any reasonable interpretation of 'subjective'.

Jaynes believes there to be an objective connection between the data available to an experimenter and his subjective Gibbs entropy. In fact, the subjective Gibbs entropy can be thought of as a property of the data rather than of the system. The objective Gibbs entropy, on the other hand, cannot depend on which data are available to the experimenter in question. It depends on the actual state of the system. If the objective Gibbs entropy is to be seen as a property of some information, it must be seen as a property of a complete description of the system. Jaynes objects to this possibility because he says that, when all possible data are available, "the notion of entopy collapses, and we are no longer talking thermodynamics". We consider Jaynes' objection to be invalid. Of course, in any real situation we only have partial information about the system and our estimate (or prediction) of the experimental entropy will depend on this partial information, but the experimental entropy itself will depend on a complete specification of the system. The crucial factor in making a good prediction is to have information about those degrees of freedom which make large contributions to the experimental entropy[4]. Whether the predictions are good or not will depend on which information has been omitted when making the prediction.

An objectivist interpretation of the proof

The motivation behind objective theories of probability is the intuition that one ought to be able to make objectively true statements about the probabilities of events. Such programmes try to cash out probability statements in terms of long run time averages or averages over ensembles of similar systems[5].

Although Jaynes calls his proof the subjective H-theorem, objectivists may reconstrue it in such a way that it fits into the objective framework. There are several points at which suitable reinterpretations are necessary.

1) The justification for the application of the canonical ensemble in stage 2 will have to be justified using another ensemble e*. The members of e* are systems resulting from relevantly similar 'repetitions' of the isolation process. The relative numbers of members with different energies will define the appropriate probability distribution. The distribution will turn out to be the canonical one because the members of e* will be a random selection from the original ensemble e appropriate to stage 1.

2) On the objectivist view S_G isn't the property of the actual system or of our state of ignorance about the actual system, rather, it is a property of an ensemble, one of whose members is the actual system. The entropy of an ensemble is a constant which does not change with time. The connection with thermodynamic entropy is that the entropy of the ensemble will correspond to the 'average' thermodynamic entropy of its members.

3) The relaxation process which occurs in the stage 4→5 transition must be understood in terms of an increase in the number of accessible states. The non-canonical distribution relaxes back into the canonical one because energy can pass freely to and fro between the heat bath and the system.

The objective interpretation of the Gibbs entropy and the whole concept of an imaginary ensemble of similar systems is rather cumbersome. Claims about the objective Gibbs analog are claims about the physical world, or at least about an ensemble of systems that could exist. Of course one can also express partial belief in such claims. The worry about the subjective analog is that it

---

[4] Naturally the inferences can only be expected to be reliable if their premises are true, e.g., one will make false predictions if one assumes certain degrees of freedom to be active at low temperatures when they are actually 'frozen out'.

[5] Which systems are to count as 'similar' will be a matter of controversy.

claims to be an expression of partial belief, not in the thermodynamic entropy, but in the true microstate of the system. According to the subjective view the thermodynamic entropy is just the expression of ignorance appropriate to somebody who only knows the macroscopic constraints on the system. The odd thing is that this particular degree of ignorance should manifest itself as a measurable heat effect. The explanation in so far as we are able to give one is that the heat effect, like the thermodynamic constraints, are macroscopic phenomena which obey certain regularities and which are largely independent of simultaneous microscopic situation. In short, thermodynamics deals with reproducible macroscopic phenomena, one of which is the heat effect we call entropy. But here we have reached an objective bottom line. The fact that the heat effect is a reproducible one will not be cashed out simply by reference to our expectations about future experiments, although that will play a part. When we say process $A \rightarrow B$ is reproducible we must be relying on the objective fact that it has been in the past. Past reproducibility doesn't of course imply future reproducibility, but one cannot rationally predict future reproducibility without some experience of past reproducibility.

References

Brillouin L.: 1962, *Science and Information Theory*, Academic Press.

Denbigh K.G. and Denbigh J.S.: 1985, *Entropy in Relation to Incomplete Knowledge*, CUP.

Denbigh K.G.: 1989, 'Note on Entropy, Disorder and Disorganization', *Brit. J. Phil. Sci.* **40**, p.323.

Grad H.: 1967, 'Levels of Description in Statistical Mechanics and Thermodynamics', *Delaware Seminar in the Foundations of Physics*, ed. M. Bunge.

Hollinger H.B. and Zenzen M.J.: 1982 'An Interpretation of Macroscopic Irreversibility within the Newtonian Framework', *Philosophy of Science* **49**, p.309.

Jaynes E.T.: 1965, 'Gibbs vs Boltzmann Entropies', *Am. J. Phys.* **33**, p.391.

Jaynes E.T.: 1985, 'Some Random Observations', *Synthese* **63**, p.115.

Jaynes E.T.: 1988, 'The Evolution of Carnot's Principle', in *Maximum-Entropy and Bayesian Methods in Science and Engineering (Vol.1)*, p.267, eds. G.J. Erickson and C.R. Smith, Kluwer.

Skyrms B.: 1985, 'Maximum Entropy Inference as a Special Case of Conditionalisation', *Synthese* **63**, p.55.